feature of any or all the claims. Moreover, the particular embodiments disclosed above are illustrative only, as the disclosed subject matter may be modified and practiced in different but equivalent manners apparent to those skilled in the art having the benefit of the teachings herein. No limitations are intended to the details of construction or design herein shown, other than as described in the claims below. It is therefore evident that the particular embodiments disclosed above may be altered or modified and all such variations are considered within the scope of the disclosed subject matter. Accordingly, the protection sought herein is as set forth in the claims below.

1. A method comprising:

generating values of thermal couplings between a plurality of layers of a three-dimensional processor stack that includes a plurality of processor cores, wherein at least one of the plurality of processor cores is implemented in each of the plurality of layers, and wherein the values of the thermal couplings indicate temperature changes in each of the plurality of layers as a function of temperature changes in each of the other layers; and

in response to a thermal event in one of the plurality of layers, selectively throttling at least one of the plurality of processor cores implemented in the plurality of layers based on the values of the thermal couplings and measures of criticality of threads executing on the plurality of processor cores.

2. The method of claim **1**, wherein selectively throttling the at least one of the plurality of processor cores comprises selectively throttling the at least one of the plurality of processor cores based on latencies between temperature changes in the plurality of layers.

3. The method of claim **1**, further comprising:

detecting the thermal event at a first processor core implemented in a first layer; and

determining whether a first thread executing on the first processor core is a critical thread based on a measure of criticality of the first thread.

4. The method of claim **3**, wherein determining whether the first thread is a critical thread comprises determining whether the first thread is a critical thread based on at least one of a criticality indicator provided by an operating system and a value of a hardware event counter associated with the first thread.

5. A method comprising:

generating values of thermal couplings between a plurality of layers of a three-dimensional processor stack,

detecting a thermal event at a first processor core implemented in a first layer; and

determining whether a first thread executing on the first processor core is a critical thread based on a measure of criticality of the first thread; and

in response to the thermal event in one of the plurality of layers, selectively throttling at least one of a plurality of processor cores implemented in the plurality of layers based on the values of the thermal couplings and measures of criticality of threads executing on the plurality of processor cores, wherein selectively throttling the at least one of the plurality of processor cores comprises selectively throttling at least one second processor core implemented in a second layer in response to the first thread being a critical thread and a second thread executing on the second processor core not being a critical thread.

6. The method of claim **5**, wherein selectively throttling the at least one second processor core comprises selectively throttling the at least one second processor core based on a thermal coupling between the first layer and the second layer.

7. The method of claim **5**, further comprising:

determining whether the thermal event is resolved by throttling the at least one second processor core; and

throttling at least one other processor core in response to determining that the thermal event has not been resolved.

8. The method of claim **1**, wherein generating the values of the thermal couplings further comprises:

iteratively executing a predetermined code on the plurality of processor cores in each of the plurality of layers; and

measuring temperatures in each of the plurality of layers in response to executing the predetermined code on the plurality of processor cores in each of the plurality of layers; and

determining the values of the thermal couplings between the plurality of layers based on the measured temperatures.

9. The method of claim **8**, wherein determining the values of the thermal couplings comprises determining values of latencies between temperature changes in each of the plurality of processor cores in each of the plurality of layers.

10. An apparatus comprising:

a three-dimensional processor stack comprising a plurality of processor cores implemented in a plurality of layers, wherein at least one of the plurality of processor cores is implemented in each of the plurality of layers; and

a controller to:

generate values of thermal couplings between the plurality of layers based on temperatures measured in the plurality of layers, wherein the values of the thermal couplings indicate temperature changes in each of the plurality of layers as a function of temperature changes in each of the other layers; and

selectively throttle at least one of the plurality of processor cores in response to detecting a thermal event, wherein the controller is to selectively throttle the at least one of a plurality of processor cores based on values of thermal couplings between the plurality of layers and based on measures of criticality of threads executing on the plurality of processor cores.

11. The apparatus of claim **10**, wherein the controller is to selectively throttle the at least one of the plurality of processor cores based on latencies between temperature changes in the plurality of layers.

12. The apparatus of claim **10**, wherein the controller is to detect the thermal event at a first processor core implemented in a first layer and determine whether a first thread executing on the first processor core is a critical thread based on a measure of criticality of the first thread.

13. The apparatus of claim **12**, wherein the controller is to determine whether the first thread is a critical thread based on at least one of a criticality indicator provided by an operating system and a value of a hardware event counter associated with the first thread.